

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application No.: 10/627,258 § Examiner: Le, Uyen T.
Filed: July 25, 2003 § Group/Art Unit: 2163
Inventor(s): § Atty. Dkt. No: 5760-12600
Paul L. Popelka and Manish M.
Agarwal

Title: Network File System
Record Lock Recovery in a
Highly Available
Environment

SECOND AMENDED APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

This Second Amended Appeal Brief is presented in response to the Notice of Non-Compliant Appeal Brief mailed August 21, 2007 (the "Notice"). The Notice objected to the Amended Appeal Brief filed July 23, 2007, asserting that the status of amendments section does not indicate that an after final amendment was entered on December 18, 2006. Appellant has removed that first sentence of the status of amendments section (that stated there were no unentered amendments after final) to clarify the section, and have further clarified that the amendment entered on December 18, 2006 was an amendment after final rejection. Appellant respectfully requests that this appeal be considered by the Board of Patent Appeals and Interferences.

I. REAL PARTY IN INTEREST

The present application is owned by Veritas Operating Corporation, which is owned by Symantec Corp. An assignment of the present application to the owner is recorded at Reel 014364, Frame 0945.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences known to Appellant.

III. STATUS OF CLAIMS

Claims 1-37 are pending. Claims 3-37 are rejected under 35 U.S.C. § 112, second paragraph. Claims 1-2 are rejected under 35 U.S.C. § 102(a). It is these rejections that are being appealed. A copy of claims 1-37 is included in the Claims Appendix attached hereto.

IV. STATUS OF AMENDMENTS

The amendment filed on December 18, 2006 (after final rejection) was entered, and overcame the rejection under 35 U.S.C. § 101. See the Advisory Action mailed January 5, 2007.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a cluster (Fig. 2) comprising a plurality of filesystems (16A-16C) and a plurality of nodes (10A-10N) (See, e.g., specification, page 8, lines 10-22). Each of the plurality of filesystems is included in one of a plurality of service groups (1A-1C), and each of the plurality of service groups includes one or more filesystems of the plurality of filesystems (See e.g., specification, page 6, lines 22-27) Each of the plurality of nodes comprises a computer system, wherein the plurality of nodes are configured to: (i) act as a server for the plurality of filesystems (See, e.g.,

specification page 9, line 7-page 10, line 20); (ii) provide record locking services in the plurality of filesystems (See, e.g., specification page 9, line 7-page 10, line 20); and (iii) maintain a plurality of client lists (28A-28C), each of the plurality of client lists included in a respective service group of the plurality of service groups and identifying clients (14A-14M) having at least one lock on one of the one or more filesystems included in the respective service group. (See, e.g., specification page 9, line 7-page 10, line 20) (See also, e.g., Figs. 1-2 and specification page 7, line 5-page 8, line 2; page 8, line 10-page 9, line 25).

Independent claim 11 is directed to a method. The method comprises maintaining a plurality of client lists (28A-28C), each of the plurality of client lists included in a respective service group of a plurality of service groups (1A-1C) and identifying clients having at least one lock on at least one filesystem (16A-16D) included in the respective service group (1A-1C) (block 50, Fig. 4 and page 9, line 27-page 10, line 1). The method further comprises failing over a first service group of the plurality of service groups from a first node of a plurality of nodes (10A-10N) to a second node of the plurality of nodes (block 52, Fig. 4 and page 4, lines 18-23; page 10, lines 4-20; and page 15, lines 4-24). The method still further comprises the second node initiating lock recovery for locks on one or more filesystems including in the first service group responsive to the fail over using a first client list of the plurality of client lists, wherein the first client list included in the first service group (blocks 56, 58, 60, Fig. 4 and page 4, lines 20-23; page 9, lines 14-17; page 10, lines 4-20; and page 15, line 15-page 16, line 10). (See also, e.g., Figs. 1-2 and 4 and specification page 7, line 5-page 8, line 2; page 8, line 10-page 9, line 25; page 14, line 25-page 16, line 9).

Independent claim 20 is directed to a computer accessible medium (150) encoded with a plurality of instructions which, when executed in a first node of a plurality of nodes (10A-10N) in response to a fail over of a first service group of a plurality of service groups (1A-1C) from a second node of the plurality of nodes and each of the plurality of service groups including at least one filesystem (16A-16D) (block 52, Fig. 4 and Fig. 10, and page 4, lines 18-23; page 10, lines 4-20; page 14, lines 27-29; and page

15, lines 4-24), initiate lock recovery for locks on each filesystem in the first service group using a first client list of a plurality of client lists (28A-28C) (blocks 56, 58, 60, Fig. 4 and page 4, lines 20-23; page 9, lines 14-17; page 10, lines 4-20; and page 15, line 15-page 16, line 10), wherein each of the plurality of client lists is included in a respective service group of the plurality of service groups and identifies clients having at least one lock in at least one filesystem in the respective service group, and wherein the first client list is included in the first service group (See, e.g. page 7, lines 19-27). (See also, e.g., Figs. 1-2 and 4 and specification page 7, line 5-page 8, line 2; page 8, line 10-page 9, line 25; page 14, line 25-page 16, line 9).

Independent claim 27 is directed to a computer accessible medium (150) encoded with a plurality of instructions which, when executed in a first node of a plurality of nodes (10A-10N) in response to a fail over of a first service group of a plurality of service groups (1A-1C) from a second node of the plurality of nodes and each of the plurality of service groups comprises at least one filesystem (16A-16D) (block 52, Fig. 4 and Fig. 10, and page 4, lines 18-23; page 10, lines 4-20; page 14, lines 27-29; and page 15, lines 4-24), initiate lock recovery for locks on the at least one filesystem in the first service group (blocks 56, 58, 60, Fig. 4 and page 4, lines 20-23; page 9, lines 14-17; page 10, lines 4-20; and page 15, line 15-page 16, line 10), and wherein locks are maintained in the at least one filesystem in a second service group (1B, Fig. 3) of the plurality of service groups during a time period that locks in the first service group are recovered (page 7, line 18-page 8, line 3; and page 9, lines 14-26). (See also, e.g., Figs. 1-4 and specification page 7, line 5-page 8, line 2; page 8, line 10-page 9, line 25; page 14, line 12-page 16, line 9).

Independent claim 33 is directed to a method. The method comprises initiating lock recovery for locks on one or more filesystems (16A-16D) in a first service group of a plurality of service groups (1A-1C) in response to failing over the first service group (block 52, Fig. 4 and Fig. 10, and page 4, lines 18-23; page 10, lines 4-20; page 14, lines 27-29; and page 15, lines 4-24) to a first node of a plurality of nodes (10A-10N) from a second node of the plurality of nodes (blocks 56, 58, 60, Fig. 4 and page 4, lines 20-23;

page 9, lines 14-17; page 10, lines 4-20; and page 15, line 15-page 16, line 10). The method further comprises maintaining locks on one or more filesystems in a second service group of the plurality of service groups during a time period that locks in the first service group are recovered (1B, Fig. 3) (page 7, line 18-page 8, line 3; and page 9, lines 14-26). (See also, e.g., Figs. 1-4 and specification page 7, line 5-page 8, line 2; page 8, line 10-page 9, line 25; page 14, line 12-page 16, line 9).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 3-37 are rejected under 35 U.S.C. § 112, second paragraph.
2. Claims 1-2 are rejected under 35 U.S.C. § 102(a) as being anticipated by the related art section of the present application.

VII. ARGUMENT

First Ground of Rejection:

Claims 3-37 are rejected under 35 U.S.C. § 112, second paragraph. Appellant traverses this rejection for at least the following reasons.

Claims 3-10:

Appellant respectfully submits that each of claims 3-10 meet the requirements of section 112, second paragraph. With regard to claim 3, the Final Office Action mailed October 31, 2006 ("Office Action") alleges that the "failing over" language is ambiguous, asserting that it is not clear what fails over to what. Appellant respectfully disagrees.

Claim 3 recites "the plurality of nodes are configured to fail over a first service group of the plurality of service groups from a first node of the plurality of nodes to a second node of the plurality of nodes." Thus, the service group clearly fails over from one node to another. There is no other way to read the above language, and thus the language is not even the least bit ambiguous.

For at least the above stated reasons, Appellant submits that the rejection of claim 3 is in error and requests reversal of the rejection. Claims 4-10 depend from claim 3, and are only rejected due to their dependence on claim 3. Accordingly, the rejection of claims 4-10 is also in error, and Appellant requests reversal of the rejection.

Claims 11-19:

Appellant respectfully submits that each of claims 11-19 meet the requirements of section 112, second paragraph. With regard to claim 11, the Office Action asserts that the nodes do not seem to be related to the service group. Appellant respectfully

disagrees. Claim 11 recites "failing over a first service group of the plurality of service groups from a first node of a plurality of nodes to a second node of the plurality of nodes". It is clear that a service group is an entity managed by a node. Specifically, for example, a first service group is failed over from a first node to a second node. There is nothing even remotely unclear about this recitation, nor is anything unclear about the relationship of nodes and service groups.

Also, the Office Action alleges that the failing over language is ambiguous, asserting that it is not clear what fails over to what. Appellant respectfully disagrees. As highlighted above, the service group fails over from one node to another. There is no other way to read the above language, and thus the language is not even the least bit ambiguous.

The Office Action also asserts that the language of "configured to" does not require any actual function to be performed by any node. Appellant respectfully disagrees. The definition of "configure" (from Merriam Webster's Collegiate Dictionary, Tenth Edition) is "to set up for operation esp. in a particularly way". Thus, "configure to" requires that the subject of the sentence (the nodes, in this case) have a structural and functional set up to realize the various features recited in claim 11. Clearly, then, a system that does not prohibit the features does not meet the features listed after "configured to", as asserted in the Office Action. Only a system that is expressly defined as having such a configuration would meet the features.

For at least the above stated reasons, Appellant submits that the rejection of claim 11 is in error and requests reversal of the rejection. Claims 12-19 depend from claim 11, and are only rejected due to their dependence on claim 11. Accordingly, the rejection of claims 12-19 is also in error, and Appellant requests reversal of the rejection.

Claims 20-37:

Appellant respectfully submits that each of claims 20-37 meet the requirements of section 112, second paragraph. With respect to claims 20, 27, and 33, the Office Action

repeats the assertion that it is not clear how the filesystems, nodes, and service groups are related and what is failing over to what. With respect to the relationship of the nodes to service groups and what is failing over to what, Appellant respectfully submits that the above remarks with regard to claims 3 and 11 illustrate why these claims are clear in this regard. As for the relationship of the filesystem to other claim features, Appellant notes that, for example, claim 20 recites "each of the plurality of service groups including at least one filesystem". Accordingly, the relationship of service group and filesystem is clear (a service group includes at least one file system). Since the relationship of service group and node is clear as highlighted above (the service group is an entity managed by a node), the relationship of filesystem and node is also clear through the relationship of service group and node. That is, a filesystem is included in a service group, which is an entity managed by a node.

For at least the above stated reasons, Appellant submits that the rejection of claims 20, 27, and 33 is in error and requests reversal of the rejection. Claims 21-26 (dependent from claim 20), 28-32 (dependent from claim 27), and 34-37 (dependent from claim 33) are only rejected due to their dependence on their respective independent claims. Accordingly, the rejection of claims 21-26, 28-32, and 34-37 is also in error, and Appellant requests reversals of the rejection.

Second Ground of Rejection:

Claims 1-2 are rejected under 35 U.S.C. § 102(a) as being anticipated by the related art section of the present application. Appellant respectfully traverses for at least the following reasons.

Claim 1:

For a reference to anticipate a claim, that reference must teach EACH and EVERY feature of the claim. The related art section of the present application does not teach at least "each of the plurality of client lists included in a respective service group of the plurality of service groups and identifying clients having at least one lock on one of the one or more filesystems included in the respective service group" as recited in claim 1.

The Office Action asserts that the related art section of the present application teaches the above features, stating that the claimed cluster is met by any network filesystem of the related art. Appellant respectfully disagrees. The network filesystems in the related art section make no mention of service groups. Furthermore, the Office Action alleges that the cluster of the related art clearly has nodes configured to perform (i), (ii), and (iii) since it is part of a network and allows lock recovery using a client list. The related art section does describe a client list. However, the related art teaches: "the server also maintains a list, in nonvolatile memory, of which clients have locks in any of the filesystems served by that server. If the server crashes or reboots, the client list is read after the server is brought back up and the clients are notified of the server crash/reboot." (specification, page 2, lines 15-18). Thus, the client list is maintained on a per-server basis, and can only be used when that particular server is rebooted.

The Office Action responds to the above argument by asserting that claim 1 does not require that the client list be useable when a particular server is not rebooted. Appellant does not disagree with this assertion. However, since the relevant art section teaches that each individual server has its own client list in its own non-volatile memory, the server of the relevant art section cannot teach or suggest the plurality of nodes (each

of which is a computer system, as recited in claim 1) are configured to: (i) act as a server for the plurality of filesystems; (ii) provide record locking services in the plurality of filesystems; and (iii) maintain a plurality of client lists. Each individual server of the related art section operates individually, not collectively as recited in claim 1.

Furthermore, the above assertion of the Office Action does nothing to overcome the fact that the related art section does not teach or suggest filesystems included in service groups, nor the fact that the related art section does not teach service groups at all.

The related art section goes on to describe a cluster in an HA setting (specification page 2, lines 22-28), but this has nothing to do with the fileserver discussion that precedes it. Furthermore, the last paragraph of the related art describes why the related art network filesystems and the related art HA clusters cannot be combined successfully (specification, page 3, lines 1-20).

For at least all of the above stated reasons, Appellant submits that the section 102 rejection of claim 1 is in error and request that the rejection be rescinded.

Claim 2:

Claim 2 depends from claim 1, and the rejection of claim 2 is in error for at least the above stated reasons as well. Furthermore, claim 2 recites a combination of features including: "each of the plurality of client lists is stored in at least one of the one or more filesystems included in the respective service group". The Office Action asserts that page 2 of the related art section teaches the above highlighted features. Appellant respectfully disagrees. Page 2 teaches that the client list is stored on the server, in non-volatile memory (See specification, page 2, lines 15-18). This does not anticipate storing the client list in the filesystem.

For at least all of the above stated reasons, Appellant submits that the section 102 rejection of claim 2 is in error and request that the rejection be rescinded.

VIII. CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejections of claims 1-37 are erroneous, and reversal of the decision is respectfully requested.

The Commissioner is authorized to charge the appeal brief fee of \$500 and any other fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5760-12600/LJM. This Appeal Brief is submitted with a return receipt postcard.

Respectfully submitted,

/Lawrence J. Merkel/
Lawrence J. Merkel, Reg. No. 41,191
Agent for Appellant

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8850

Date: August 30, 2007

IX. CLAIMS APPENDIX

The claims on appeal are as follows.

1. A cluster comprising:

a plurality of filesystems, each of the plurality of filesystems included in one of a plurality of service groups, and each of the plurality of service groups including one or more filesystems of the plurality of filesystems; and

a plurality of nodes, each of the plurality of nodes comprising a computer system, wherein the plurality of nodes are configured to: (i) act as a server for the plurality of filesystems; (ii) provide record locking services in the plurality of filesystems; and (iii) maintain a plurality of client lists, each of the plurality of client lists included in a respective service group of the plurality of service groups and identifying clients having at least one lock on one of the one or more filesystems included in the respective service group.

2. The cluster as recited in claim 1 wherein each of the plurality of client lists is stored in at least one of the one or more filesystems included in the respective service group.

3. The cluster as recited in claim 1 wherein the plurality of nodes are configured to fail over a first service group of the plurality of service groups from a first node of the plurality of nodes to a second node of the plurality of nodes, and wherein the second node is configured to initiate lock recovery for locks in the one or more filesystems included in the first service group responsive to the fail over.

4. The cluster as recited in claim 3 wherein the second node is configured to read a first client list of the plurality of client lists, the first client list included in the first service group, to initiate lock recovery.

5. The cluster as recited in claim 3 wherein the second node is configured to initiate lock recovery by notifying each of the clients in the first client list that the clients should reclaim locks previously granted on a filesystem within the first service group.
6. The cluster as recited in claim 5 wherein the second node is configured to notify each of the clients by transmitting one or more server identifiers within the first service group to the clients.
7. The cluster as recited in claim 6 wherein each of the one or more server identifiers comprises an internet protocol address.
8. The cluster as recited in claim 5 wherein the second node is configured to initiate a period of time for the clients in the first client list to reclaim locks, wherein the second node is configured not to grant new locks on filesystems within the first service group during the period.
9. The cluster as recited in claim 8 wherein, if the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the second node is configured to grant a lock requested by a client in a filesystem within the second service group during the period.
10. The cluster as recited in claim 3 wherein, if the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the second node is configured to not initiate lock recovery for locks in the second service group responsive to the fail over of the first service group.
11. A method comprising:

maintaining a plurality of client lists, each of the plurality of client lists included in a respective service group of a plurality of service groups and identifying clients having at least one lock on at least one filesystem

- included in the respective service group;
- failing over a first service group of the plurality of service groups from a first node of a plurality of nodes to a second node of the plurality of nodes; and
- the second node initiating lock recovery for locks on one or more filesystems including in the first service group responsive to the fail over using a first client list of the plurality of client lists, the first client list included in the first service group.
12. The method as recited in claim 11 further comprising storing each of the plurality of client lists in at least one of the one or more filesystems included in the respective service group.
13. The method as recited in claim 11 wherein initiating lock recovery comprises notifying each of the clients in the first client list that the clients should reclaim locks previously granted on at least one filesystem in the first service group.
14. The method as recited in claim 13 wherein notifying each of the clients comprises transmitting one or more server identifiers included in the first service group.
15. The method as recited in claim 13 further comprising:
- initiating a period of time for the clients in the first client list to reclaim locks in the filesystems included in the first service group; and
- the second node not granting new locks in the filesystems included in the first service group during the period.
16. The method as recited in claim 15 wherein the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups,

the method further comprising the second node not interrupting locking services for the second service group during the period.

17. The method as recited in claim 16 wherein not interrupting locking services comprises granting a lock in a filesystem within the second service group during the period.
18. The method as recited in claim 15 wherein the first node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the method further comprising the first node not interrupting locking services for the second service group during the period.
19. The method as recited in claim 11 wherein the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the method further comprising the second node not interrupting locking services for the second service group during the period.
20. A computer accessible medium encoded with a plurality of instructions which, when executed in a first node of a plurality of nodes in response to a fail over of a first service group of a plurality of service groups from a second node of the plurality of nodes and each of the plurality of service groups including at least one filesystem, initiate lock recovery for locks on each filesystem in the first service group using a first client list of a plurality of client lists, wherein each of the plurality of client lists is included in a respective service group of the plurality of service groups and identifies clients having at least one lock in at least one filesystem in the respective service group, and wherein the first client list is included in the first service group.
21. The computer accessible medium as recited in claim 20 wherein each of the plurality of client lists is stored in the respective service group.
22. The computer accessible medium as recited in claim 20 wherein the plurality of

instructions, when executed, initiate lock recovery by notifying each of the clients in the first client list that the clients should reclaim locks previously granted on a filesystem in the first service group.

23. The computer accessible medium as recited in claim 22 wherein notifying each of the clients in the first client lists comprises transmitting one or more server identifiers included in the first service group.

24. The computer accessible medium as recited in claim 22 wherein the plurality of instructions, when executed:

initiate a period of time for the clients in the first client list to reclaim locks on the at least one filesystem in the first service group; and

do not grant new locks on the at least one filesystem in the first service group during the period.

25. The computer accessible medium as recited in claim 24 wherein, if the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the plurality of instructions, when executed, grant a lock requested by a client on a filesystem in the second service group during the period.

26. The computer accessible medium as recited in claim 20 wherein, if the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the plurality of instructions, when executed, do not initiate lock recovery for locks on filesystems in the second service group responsive to the fail over of the first service group.

27. A computer accessible medium encoded with a plurality of instructions which, when executed in a first node of a plurality of nodes in response to a fail over of a first service group of a plurality of service groups from a second node of the plurality of nodes and

each of the plurality of service groups comprises at least one filesystem, initiate lock recovery for locks on the at least one filesystem in the first service group, and wherein locks are maintained in the at least one filesystem in a second service group of the plurality of service groups during a time period that locks in the first service group are recovered.

28. The computer accessible medium as recited in claim 27 wherein the first node is acting as a server for the at least one filesystem in the second service group.
29. The computer accessible medium as recited in claim 27 wherein a different node of the plurality of nodes is acting as the server for the at least one filesystem in the second service group.
30. The computer accessible medium as recited in claim 27 wherein the plurality of instructions, when executed:
 - initiate a period of time for clients to reclaim locks on the at least one filesystem in the first service group; and
 - do not grant new locks on the at least one filesystem in the first service group during the period.
31. The computer accessible medium as recited in claim 30 wherein the plurality of instructions, when executed, grant a lock requested by a client on one of the at least one filesystems included in the second service group during the period.
32. The computer accessible medium as recited in claim 27 wherein the plurality of instructions, when executed, do not interrupt locking services for the second service group.
33. A method comprising:

initiating lock recovery for locks on one or more filesystems in a first service group of a plurality of service groups in response to failing over the first service group to a first node of a plurality of nodes from a second node of the plurality of nodes; and

maintaining locks on one or more filesystems in a second service group of the plurality of service groups during a time period that locks in the first service group are recovered.

34. The method as recited in claim 33 wherein the first node is acting as a server for the one or more filesystems in the second service group.

35. The method as recited in claim 33 wherein a different node of the plurality of nodes is acting as the server for the one or more filesystems in the second service group.

36. The method as recited in claim 33 further comprising:

initiating a period of time for clients to reclaim locks on the one or more filesystems in the first service group; and

not granting new locks on the one or more filesystems in the first service group during the period.

37. The method as recited in claim 36 further comprising granting a lock requested by a client on a filesystem in the second service group during the period.

X. EVIDENCE APPENDIX

No evidence submitted under 37 CFR §§ 1.130, 1.131 or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

XI. RELATED PROCEEDINGS APPENDIX

There are no related proceedings known to Appellant.